

Building a Data System to Produce, Archive and Distribute Global Data Products from MODIS

Ed Masuoka/NASA¹, Curt Tilmes/Raytheon-STX, Carey Noll/NASA, Dr. Shahin Samadi/Bureau of National Affairs, and Nick Ruggiero, Dr. William Englemeyer and Dr. Jack Schols of SAIC/GSC

¹Goddard Space Flight Center, Code 922, Greenbelt, MD 20771
(301) 614-5515 voice, (301) 614-5269 fax, Email: emasuoka@pop900.gsfc.nasa.gov

Abstract – An emergency backup processing system for NASA's EOSDIS was developed to produce, archive and distribute global science products from the MODerate resolution Imaging Spectroradiometer (MODIS) instrument which will be launched on the EOS AM-1 satellite. This system will produce up to 400GB of MODIS products per day and distribute up to 500GB of data via the Internet and on DLT and Exabyte tapes.

INTRODUCTION

The MODerate resolution Imaging Spectroradiometer (MODIS) will be launched on a series of Earth Observing System (EOS) satellites starting in 1998 on EOS AM-1. With 36 spectral bands covering the visible, near wave and short wave infrared, MODIS produces over 40 global science data products including sea surface temperature, ocean color, cloud properties, vegetation indices and land cover. To produce and store these data products requires computers which can perform 6 billion floating point operations per second continuously and an archive which will store more than 200 trillion bytes of data per year. The data processing load and data product volume from MODIS drove the performance requirements of the science data processing system for EOS (EOSDIS) which will produce, archive and distribute the data products for EOS instruments.

In March 1997, the MODIS Team was selected to develop a computing system to support early science studies, data product validation and to serve as a backup processing system for the EOSDIS. This system had to: be fully operational by the EOS AM-1 launch scheduled for June 1998; meet processing and data distribution requirements for the first 6 months after the launch and run MODIS science software written for the EOSDIS. This paper describes: the software development process for the MODIS Emergency Backup System (MEBS), its hardware and software architecture and how the system architecture has evolved over time to reach the desired level of processing and archive performance and address new user requirements.

SOFTWARE DEVELOPMENT PROCESS

In developing MEBS, we were aware from the beginning of

the pressure of meeting schedule. It was essential that the MEBS system be ready to produce, archive and distribute MODIS products by the launch of EOS AM-1. Despite the need to develop MEBS quickly to meet the launch date, an important goal of the software team was to design a system which would be capable of supporting development of improved algorithms and quality assurance of data products over the 15 year life of MODIS sensors.

MEBS software was developed following the Evolutionary Delivery approach [1, 2]. In Evolutionary Delivery after the initial requirements analysis and the design of the core system architecture are completed, the system is developed through a series of incremental deliveries. After each delivery, customer feedback is sought and their suggestions are incorporated in subsequent deliveries. We chose this approach to system development because: we felt we understood the core functionality of the MEBS system; we had several software packages which we planned to modify and combine to provide basic functionality; and we needed user feedback and opportunities to adjust the design in areas where requirements were not well defined.

During the development of the MEBS system, software releases were delivered every 1 to 3 months. The first prototype of the job scheduler in February 1997 was followed by a series of deliveries from May through December, which included a data ordering interface, data product archive, support for automated tape libraries and MODIS science. In February 1998 the capability to produce all weekly science products and to run unattended for 3 days was delivered and demonstrated to the MODIS Science Team. Each delivery of MEBS has resulted in requests for new capabilities and/or modifications of existing features. These requests are reviewed by a configuration control board, which decides if a requested feature should be implemented and if so in what release.

MEBS SOFTWARE

The MEBS system is composed of 5 subsystems, Data Ingest, Product Generation, Archiver, MEBDOS and Monitor Production, which interact via messages and three data stores (Product Storage, Archive Request Queue and Product Catalog) as shown in Fig. 1.

In designing MEBS, we recognized the need to use existing software from other satellite processing systems and employ commercial software where possible to save time in development and software maintenance. Three science data processing systems, SeaWiFS, TSDIS and EOSDIS, were examined to determine if portions of their software might be modified to support MODIS data production. The SeaWiFS system was selected because it had undergone extensive testing, produced an ocean color product suite similar to that needed for MODIS and the SeaWiFS development team offered to assist in modifying their software to support MODIS processing. The SeaWiFS production system [3] is capable of scheduling and tracking hundreds of jobs on over 40 processors via a graphical operations interface.

The Data Ingest subsystem handles raw instrument data sets from the LZAARDS (Level Zero Active Archive Data System) and ancillary data sets from the Goddard Distributed Active Archive Center. When these data sets arrive in their ftp directories on the MEBS system, metadata describing these data sets are stored in the Product Catalog and the Product Generation component releases those jobs which were awaiting the arrival of these data sets.

The Product Generation subsystem consists of 4 components which automate one or more days of MODIS data production. The first, Scheduling, assists the operations manager in developing a production schedule of jobs for one or more days. The jobs are then stored in a production parameters database. The Loader reads the production parameter database, generates requests to stage the required input for each job and provides information about each job to the Scheduler which oversees the execution of all product generation processes. At the conclusion of a job products are stored in on-line storage where the Extractor retrieves metadata from each product and updates the Product Catalog. Products will remain in on-line storage until all jobs that require them have run or the amount of free disk space drops below a predetermined amount.

The Archiver subsystem of MEBS consists of 4 components which manage on-line disk space and store/retrieve files from the automated tape libraries. A Disk Space Monitor component compares the available on-line disk space with the lowest value of free space allowed. If free space falls below this value, then files which have been saved in the tape archive are deleted from disk. The Archive Controller component monitors the Archive Request queue and schedules transfers of products to and from the tape libraries. The Archive Store/Restore component, which handles the transfers to and from tape, was written so that the commercial software package used to manage the tape libraries can be easily swapped out. The last component is a COTS software package which handles file backup and restore

with the Ampex tape libraries. To date we have used Silicon Graphics IRIX Networker but are planning to switch to Network Solutions QuickRestore, which provides better performance during file restores from the tape libraries.

The MEBDOS subsystem utilizes a World Wide Web interface, which enables a scientist to order data sets by product name, geographic and temporal coverage as illustrated in Fig. 2. The map of the world and selection box are implemented as JAVA applets to enable the user to select data products by dragging and resizing the selection box as well as by typing in latitudes and longitudes of the corners of the selection box. It also provides links to a description of each data product and displays thumb nail images and browse products to assist in ordering. This component also indicates whether the requested data are in on-line disk storage and if not will submit a request to the Archiver subsystem to retrieve the data product(s). If the requester has asked for the data on 8mm Exabyte tape, DLT tape or via an ftp push to disk space on their site, a message confirming shipment will be sent to their email address. If the requester has asked that the data be staged or retrieval by ftp pull they will be notified when the data are on-line and have 4 days to retrieve their data sets. The Data Ordering subsystem also handles user registration and validation to insure that only approved users are able to order products from the system.

The Monitor Production subsystem receives logs, statistics and alerts from all other MEBS subsystem and generates reports on daily, weekly and monthly production, status of system hardware, free space in on-line disk storage and in the tape archive and volumes of products ordered and distributed. It is the subsystem which is expected to change the most over time as we extend the reporting to include tracking production in Distributed Active Archive Centers and MODIS Science Computing Facilities producing MODIS products.

MEBS HARDWARE

In the current MEBS architecture, processing occurs on 3 Silicon Graphics Inc. (SGI) Origin 2000 each with 16 R10000 processors connected by Fibre Channel and FDDI networks to workstations which support scientists engaged in algorithm development, product validation and daily quality assurance of data products. On-line data storage for science data processing and data distribution is provided by 1,800GB of Fibre Channel attached RAID. SGI Challenge DM workstations serve as tape hosts for two 12.5TB tape libraries which provide a near-line archive of science data products. These Challenge DM servers also run job scheduling software and the data ordering system.

During our testing, science software read and wrote data products in HDF format to locally attached RAID at 4MB to

12 MB/second resulting in data rates of 64 to 192MB/second per system. After data products were produced, they were copied to working storage for the tape libraries at a rate of 50MB/second and then written to tape at rate of 30MB/second. At this rate, 600GB of MODIS daily products can be written to tape in under 3 hours though in actual operations tape staging and destaging will occur in parallel with science processing and in general all tape drives will not be operating at once to copy a day's worth of data. At present, the processing efficiency of the science algorithms and the speed of the Challenge DM workstation as tape host and job scheduler are factors which limit system throughput.

In May 1998 we are moving to the hardware architecture shown in Fig. 3. The 3 SGI Origin 2000 systems will be linked to form a single processing system with 48 R10000 processors and the 2 Ampex 812 tape libraries and 1.8TB of Fibre Channel RAID disks will be connected to this system. The new architecture will enable us to manage a single pool of disk space rather than separate file systems on each Origin 2000 and thereby simplify the management of free space on disks. It will also free up bandwidth on the Fibre Channel Loop providing better performance when we push data to Fibre Channel attached SGI O2 workstations and a Power Challenge compute server involved in quality assurance of MODIS land data products. Most importantly it will remove a large processing load from the SGI Challenge DM server which was handling all tape I/O and data transfers from the Origin 2000 based production system as well as job scheduling. This is essential as the Challenge DM system was constantly at 100% cpu utilization during tests with a single Origin 2000. In the new architecture it will handle job scheduling and run the Sybase database which stores job scheduling tables, the Product Catalog and production monitoring statistics while the data transfers to the Ampex libraries will go through the Origin 2000, which is better equipped to handle the 15MB/second I/O rate of each of the 6 Ampex tape drives.

MODIS DATA PRODUCTS

The MODIS instrument images the Earth in 36 spectral bands from 0.4 μm to 14.5 μm at spatial resolutions of 250m to 1km. With an average data rate of 6Mb/second the MODIS instrument collects 70GB of raw data per day from which over 40 global science data products will be produced. A number of products are available at several different temporal resolutions (daily, weekly, ... yearly), spatial resolutions and/or in different projections (equal angle vs. equal area) which brings the total number of MODIS data sets which can be ordered to over 70.

Table 1 summarizes the data products that will be produced by MEBS in the first 6 months following launch and gives

the temporal coverage of the product, the size of one instance of the product and the daily volume for each product. The column labeled "Level" refers to the level of processing performed to produce the data product. The Level 0 product is raw instrument data. Level 1 products include: uncalibrated instrument data which are unpacked and placed into 16bit integers, geolocation fields that provide the earth location for each 1km spatial element in a scene and a calibrated radiances product. The Level 2 products contain geophysical parameters, such as Snow Cover or Vegetation Indices, in the original instrument scan geometry, i.e. the spatial elements in each scanline are not earth located. While Level 1 and Level 2 products make up the largest portion of the daily MODIS data volume, the majority of the MODIS products are Level 3 or Level 4 products. To make a Level 3 product, the appropriate Level 2 products and the geolocation field product are read and the Level 2 products are combined to make a global earth located data set for the specific geophysical parameter. Level 3 products are available in at daily, weekly, monthly, quarterly and yearly versions and have observations on global grids whose cell size range from 250m to 1km for land products and up to 111km for Climate Modeling Grid products which are used as inputs in global circulation models. Level 1 and Level 2 data products contain or were produced from a 5 minute swath of data. Each 5 minute data set covers an area 2,330km by 2,000km and a single day's worth of a product consists of either 144 data sets for day-only products, like vegetation indices, or 288 data sets for products made for both day and night, like sea surface temperature. The relatively large size of the Level 2 products and the fact that they are not resampled to a fixed global grid will limit their utility for most users. Level 3 and Level 4 products which are produced on a family of nested global grids will be the choice of most researchers [4-6]. Samples of most MODIS products in their at-launch data formats can be obtained from the MEBS WWW home page at <http://ftpwww.gsfc.nasa.gov/MODIS/SDST/mebs/>. Fig. 4 is a calibrated MODIS Band 19 (.92 μm) which was produced in MEBS from synthetic data that has been earth located by applying the corrections from the geolocation fields product.

CONCLUSIONS

The MODIS Emergency Backup System is ready to produce, archive and distribute products for the MODIS instrument after the launch of the EOS AM-1 satellite in 1998. It has passed an extensive production and archiving test in which a week's worth of MODIS products were produced. Future directions for the MEBS development include additional documentation of the software and its capabilities and porting the system to a desktop workstation to facilitate algorithm development and testing for individual scientists.

ACKNOWLEDGMENTS

The authors thank Dr. Gene Feldman, Dr. Mary Cleave and John Wilding of the SeaWiFS Project for their help in adapting the SeaWiFS production system for use in the MEBS and Rick Obenschain and Dr. H. K. Ramapriyan of the ESDIS Project for their funding of the system development. Dr. Robert Murphy, Dr. Bruce Guenther, Dr. David Roy, Dr. Stanley Scott, Dr. Chris Lynnes, Robert Wolfe and Paul Fisher critiqued the software design of MEBS in the Preliminary and Critical Design Reviews.

REFERENCES

- [1] S. McConnell, Rapid Development: Taming Wild Software Schedules, Microsoft Press, 647p., 1996.
- [2] M. Freeze, W. Englemeyer, N. Ruggiero, MEBS Software Development Plan (SDP), 1997. At URL: <http://ftpwww.gsfc.nasa.gov/MODIS/SDST/docs.html/>
- [3] SeaWiFS home page is located at on the WWW at URL: http://seawifs.gsfc.nasa.gov/SEA_WIFS.html
- [4] Esaias, W et al, "An Overview of MODIS Capabilities for Ocean Science Observations", TGARS EOS-AM Special Issue, in press.
- [5] Justice, C. et al, "The MODERate Imaging Spectroradiometer (MODIS) Mission to Planet Earth enhanced land product description", TGARS EOS-AM Special Issue, in press.
- [6] Masuoka, E., A. Fleig, R. Wolfe and F. Patt, "Key characteristics of MODIS data products", TGARS EOS-AM Special Issue, in press.

Table 1. MODIS products and their data volumes

Description	Level	Discipline	Temporal Coverage	File Size (MB)	Files Created per Day	Volume of Data Archived per Day (MB)
2 hours of Instrument Path	0	Level 1	2 hours	5850	12	70.20
Level 1A Raw Counts	1	Level 1	5 minutes	328	288	94.76
Level 1A Geolocation Fields	1	Level 1	5 minutes	67	288	16.33
Level 1B Collocated Radiance Scan	1	Level 1	5 minutes	221	288	83.50
Level 1B Collocated Radiance 250m bands	1	Level 1	5 minutes	65	288	16.72
Level 1B Collocated Radiance 500m bands	1	Level 1	5 minutes	67	288	16.42
Collocated Radiance Collocated Collocated	1	Level 1	5 minutes	18	288	6.24
20 Surface Reflectance	2	land	1 day	422	288	121.96
Level 2 Cloud and Surface Classification Maps	2	atmosphere	5 minutes	45	288	13.04
Atmospheric Profiles	2	atmosphere	5 minutes	26	288	7.46
Near IR Precipitable Water	2	atmosphere	5 minutes	8	288	2.19
Level 2 Land Aerosol Product	2	atmosphere	5 minutes	10	144	1.43
Cloud Product	2	atmosphere	5 minutes	11	288	3.30
Snow Cover	2	land	5 minutes	11	144	1.56
Sea Ice Extent	2	land	5 minutes	11	144	1.56
Ocean Color (Water Leaving Radiance)	2	ocean	5 minutes	148	144	21.48
Ocean Color (Group 2 parameters)	2	ocean	5 minutes	128	144	16.97
Ocean Color (Group 1 parameters)	2	ocean	5 minutes	100	144	22.96
Sea Surface Temperature - 2 day and 2 night products	2	ocean	5 minutes	42	288	12.06
Land Surface Temperature	2	land	5 minutes	24	288	7.00
Level 30 Thermal Anomalies	2	land	1 day	21	276	15.27
5 Day Gridded Surface Reflectance	3	land	5 days	249	288 per 5 days	8.00
Gridded BRDF/Albedo	3	land	16 day	166	288 per 16 days	3.00
16 day BRDF/Albedo, Climate Modeling Grid 25 degrees	3	land	16 day	16	1 per 16 days	0.00
16 day Gridded Vegetation Index	3	land	16 day	111	288 per 16 days	2.00
Monthly Gridded Vegetation Index	3	land	1 month	111	288 per month	1.00
16 day Vegetation Index - Climate Modeling Grid 0.5 degree	3	land	16 days	16	1 per 16 days	0.00
Monthly Vegetation Index - Climate Modeling Grid 0.5 degree	3	land	1 month	16	1 per month	0.00
5 Day Gridded Thermal Anomalies	3	land	5 days	145	288 per 5 days	6.50
Monthly Gridded Thermal Anomalies	3	land	1 month	221	288 per month	2.00
Weekly Land Surface Temperature	3	land	8 days	27	1 per 8 days	1.00
Daily Land Surface Temperature - Climate Modeling Grid 0.5 degree	3	land	1 day	4	1	0.00
Land Area Index Daily Product	3	land	1 day	28	288	6.08
Monthly Gridded Land Cover Database	3	land	1 month	443	288 per month	4.00
Quarterly Gridded Land Cover Type	3	land	3 months	16	288 per 3 months	0.06
60 day Land Cover - Climate Modeling Grid 0.5 degree	3	land	3 months	4	1 per 3 months	0.00
5 Day Land Surface Temperature - Climate Modeling Grid 0.5 degree	3	land	5 days	4	1 per 5 days	0.00
Monthly Land Surface Temperature - Climate Modeling Grid 0.5 degree	3	land	1 month	4	1 per month	0.00
Monthly Land Area Index - Climate Modeling Grid 0.5 degree	3	land	1 month	4	1 per month	0.00
Monthly BRDF/Albedo, Climate Modeling Grid 25 degrees	3	land	1 month	16	1 per month	0.00
Monthly Land Cover Change	3	land	1 month	221	288 per month	2.00
Weekly Gridded Ocean Productivity	3	ocean	8 days	1620	1 per week	0.20
Yearly Gridded Ocean Productivity	3	ocean	1 year	3280	1 per year	0.41
Daily Sea Surface Temperature - 2 day and 2 night products	3	ocean	1 day	620	4	8.48
Weekly Sea Surface Temperature - 2 day and 2 night products	3	ocean	8 days	620	4 per week	0.31
Daily Ocean Color (Water Leaving Radiance Group)	3	ocean	1 day	620	12	7.44
Daily Ocean Color (Group 2 parameters)	3	ocean	1 day	620	11	6.82
Daily Ocean Color (Group 1 parameters)	3	ocean	1 day	680	13	6.84
Weekly Ocean Color (Water Leaving Radiance Group)	3	ocean	8 days	620	12 per week	0.93
Weekly Ocean Color (Group 2 parameters)	3	ocean	8 days	620	11 per week	0.93
Weekly Ocean Color (Group 1 parameters)	3	ocean	8 days	680	13 per week	1.11
Daily Gridded Atmosphere Product, Climate Modeling Grid 1 degree	3	atmosphere	1 day	450	1	0.46
Monthly Gridded Atmosphere Product, Climate Modeling Grid 1 degree	3	atmosphere	1 month	500	1 per month	0.08
Daily Gridded Snow Cover	3	land	1 day	11	354	4.00
Daily Gridded Sea Ice	3	land	1 day	8	354	1.00
Weekly Gridded Snow Cover	3	land	8 days	23	130 per 8 days	0.20
Weekly Gridded Sea Ice Cover	3	land	8 days	2	130 per 8 days	0.20
Weekly Sea Ice Cover - Climate Modeling Grid 0.25 degree	3	land	8 days	16	1 per 8 days	0.00
Weekly Sea Ice - Climate Modeling Grid 0.25 degree	3	land	8 days	16	1 per 8 days	0.00
5 Day Land Area Index & FPAR	4	land	5 days	28	288 per 5 days	1.00
5 Day Land Area Index - Climate Modeling Grid 0.5 degree	4	land	5 days	16	1 per 5 days	0.00
5 Day Net Primary Productivity	4	land	5 days	28	288 per 5 days	1.00
Yearly Net Primary Productivity - Climate Modeling Grid 0.5 degree	4	land	1 year	28	288 per year	0.00
Monthly Land Area Index - Climate Modeling Grid 0.25 degree	4	land	1 month	16	1 per 30 days	0.00
Yearly Net Primary Productivity - Climate Modeling Grid 0.5 degree	4	land	1 year	4	1 per year	0.00
Yearly Net Primary Productivity - Climate Modeling Grid 0.5 degree	4	land	1 year	4	1 per year	0.00
Total Archived Volume (GByte day)						808.76

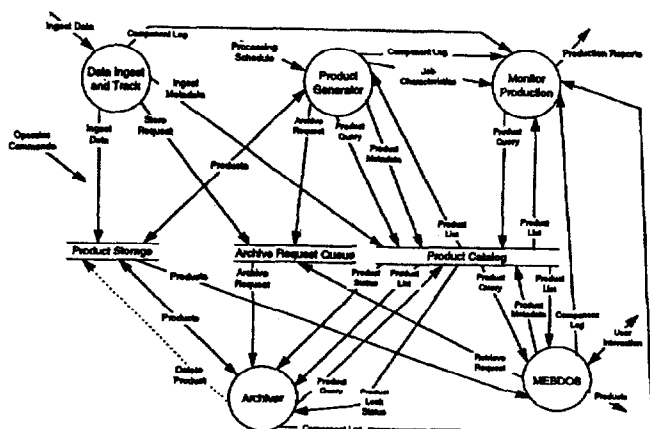


Figure 1. Top level MEBS system design


MODIS Emergency Backup Data Ordering System
Browse Selection

A list of available data granules can be generated by completing the form below.

Satellite:

Product:

The latitude/longitude locations are given by coordinates of a rectangular box:



or Defined file:

Start of Data: Date: Time:

End of Data: Date: Time:

File added to system since: Date: Time:

Day/Night: (of data acquisition)

Select data with known serious flaws:

Data version:

Figure 2. MEBS data ordering page on the WWW.




Figure 4. MODIS Level 1B calibrated radiances (Band 19) produced from synthetic test data in MEBS.

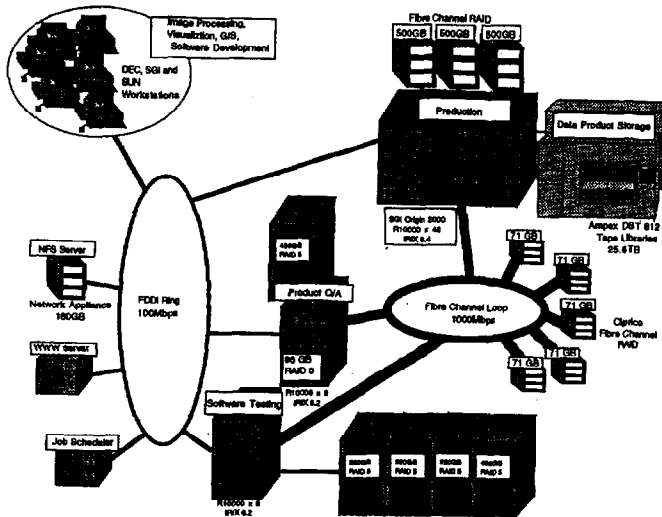


Figure 3. MEBS computing systems in the MODIS Team Leader's Science Computing Facility